arXiv:1703.09928v2 [cs.CV] 5 Apr 2017

Bundle Optimization for Multi-aspect Embedding

Qiong Zeng[1,*] Wenzheng Chen[1,*] Zhuo Han[1] Mingyi Shi[1]
Yanir Kleiman[2] Daniel Cohen-Or[2] Baoquan Chen[1] Yangyan Li[1,†]

[1]Shandong University [2]Tel Aviv University

Abstract

Understanding semantic similarity among images is the core of a wide range of computer vision applications. An important step towards this goal is to collect and learn human perceptions. Interestingly, the semantic context of images is often ambiguous as images can be perceived with emphasis on different aspects, which may be contradictory to each other. In this paper, we present a method for learning the semantic similarity among images, inferring their latent aspects and embedding them into multi-spaces corresponding to their semantic aspects. We consider the multi-embedding problem as an optimization function that evaluates the embedded distances with respect to the qualitative clustering queries. The key idea of our approach is to collect and embed qualitative measures that share the same aspects in bundles. To ensure similarity aspect sharing among multiple measures, image classification queries are presented to, and solved by users. The collected image clusters are then converted into bundles of tuples, which are fed into our bundle optimization algorithm that jointly infers the aspect similarity and multi-aspect embedding. Extensive experimental results show that our approach significantly outperforms state-of-the-art multi-embedding approaches on various datasets, and scales well for large multi-aspect similarity measures.

1. Introduction

Understanding semantic similarity among images is the core of a wide range of computer vision analysis and image retrieval applications. However, it is a particularly challenging task as it reflects how humans perceive images, a task that cannot be inferred by low-level analysis. Supervised learning is a common means of studying such semantic problem, for which, the ground truth of how humans perceive similarity among images is critical.

However, the semantic context of images is often ambiguous as images can be perceived with emphasis on different aspects (see Figure 1). One example out of many is the separation of content and style. One can claim that two images are similar due to their content and another may find two images of similar content different due to their common style. Similarities between the images may be measured in multiple aspects, which can be contradictory to each other. The key to resolve such contradictions is to disentangle similarities based on their latent aspects.

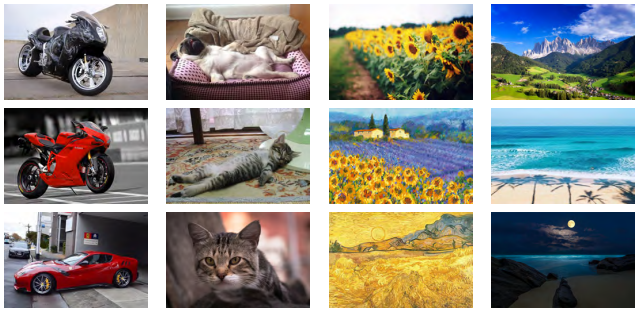

Figure 1: Often the semantic of images is ambiguous. It is unclear whether images in the second row are more similar to the top or bottom row.

Humans cannot state a meaningful quantitative measure of semantic similarity. Therefore, annotations about semantic similarity collected by crowd queries are qualitative in nature. In addition, they only contain a partial view of a whole dataset. To consolidate the partial and possibly conflicting views, the collected qualitative data is often embedded into a common Euclidean space that hosts all the images, such that the quantitative metric distances among images reflect the aggregate of the qualitative measures as much as possible [14, 21]. However, since semantic similarity may reflect various aspects, one embedding space cannot represent well multiple distances among images.

In this paper, we present a method for multi-aspect embedding, which disentangles similarity annotations based on their aspects and embeds them into multiple corresponding embedding spaces. The distances in each embedding space well represent object similarity under the corresponding aspect. The task is challenging since it encapsulates a two-fold problem. First, the similarity among images has no clear quantitative measure and thus, it must be deduced from qualitative measures. Second, the

*Joint first authors.
†Corresponding author.

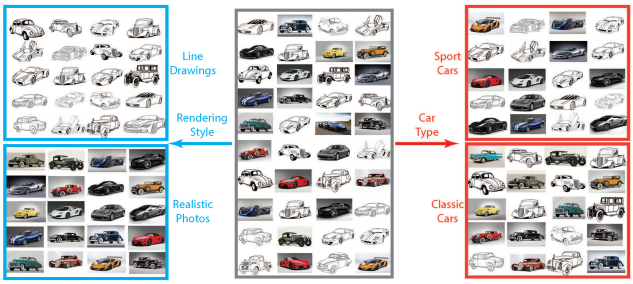

Figure 2: Different clusterings reflect similarity in different aspects, which may be contradictory to each other.

aspect that each crowd member relies on is unknown.

Our general approach to the problem is similar to the one presented by Amid and Ukkonen [3]. Qualitative queries are crowdsourced, and the answers to these qualitative queries are embedded into multiple embedding spaces. The multi-embedding is computed by optimizing an objective function that evaluates the embedded distances with respect to the qualitative queries. Each embedding represents the semantic similarities among images using one specific aspect.

A critical issue in the optimization is to infer which aspect is used in answering a particular query. Thus, each query is associated with an additional variable on top of the unknown coordinates of the embedded elements. The key idea of our approach is to collect and embed qualitative measures in *bundles*. The bundled measures necessarily share the same aspect, which significantly reduces the number of unknown variables. This directly leads to a significant increase in the accuracy of the embeddings.

More specifically, rather than collecting pairs or triples of relations one by one, the task is designed as classifying a collection of images into clusters (see Figure 2). This necessarily leads the user to use a single aspect in providing a series of qualitative measures on the collection of images. Each clustering annotation is then converted into a bundle of $T(i, j, \theta)$-like tuples, where θ indicates whether image O_i is similar to image O_j ($\theta = 1$) or not ($\theta = 0$), and fed into our bundle embedding optimization. As we shall show, the classification based task is much more cost-effective, and the optimization with tuple bundles requires less variables, leading to higher quality embeddings.

We evaluate our approach on various synthetic and crowdsourced datasets, and show that it significantly outperforms state-of-the-art multi-embedding approaches.

2. Related Work

Together with the availability of massive data, crowdsourcing allows supervised learning to be performed in a truly large scale [18], providing an efficient way to measure human perception in various contexts, such as product design [4], illustration style [9], font similarity [17] and entity matching [24]. For a recent comprehensive study of crowdsourcing, please refer to [6].

Single Embedding Metric Learning. The problem of consolidating numerous instances of information, which is often quantitative, into a single consistent space is referred to as *metric learning*, and is widely studied [10, 23, 27, 28]. Quantifying human similarity perception is challenging since it is often qualitative and relative. A number of metric learning methods focus on recovering a single embedding space from such relative similarity measures, in the form of paired comparisons [2] or relative triplets [21]. Kleiman *et al*. [14] proposed clustering queries which provide more information compared to pairs or triplets of images.

Such methods assume that similarity between two objects can be depicted by a single scalar value, thus a single embedding space can capture similarity among a set of objects. Similarity measures, which might be from different aspects, are "fused" into one embedding space. Instead, we model similarity between two objects as a multi-dimension vector, *i.e*., two objects may have different degree of similarity under different aspects. We propose to "disentangle" similarities and embed them in multiple embeddings by their aspects, which can be separately explored.

Multi-aspect Embedding. Learning multiple embeddings in general cases has not been explored much, even though it is often essential for various human-computational applications. Recent research in natural language processing has proposed a number of models in which words are associated with several corresponding embeddings based on human word similarity judgments [15, 16, 26]. However, these models use additional information such as local co-occurrence and sentence context which are not available in the general case.

Perhaps most similar to our work is the recent work by Amid and Ukkonen [3]. Their multi-view triplet embedding algorithm (MVTE) aims to reveal multiple embeddings by maximizing the sum of log-probabilities of triplet-embedding mapping over all triplets. The triplets are collected individually and it is unknown whether a pair of triplets are induced by the same aspect or different aspects. The triplet-embedding mapping is defined as an heuristic indicator function of the embedding based on distribution assumptions of the underlying embedding spaces. The method alternates between optimizing the embeddings with fixed indicators and deriving the indicators from embeddings.

Unlike MVTE, we collect image similarity annotations in the form of bundles. Each bundle contains information regarding many images, that share a single unknown

aspect. We do not make any distribution assumption of the underlying embedding spaces. Instead, we introduce a set of *aspect inference variables* that represent the mapping probabilities. These aspect inference variables are optimized simultaneously with the embedding variables, to model complex bundle-embedding relationships. In Section 6, we present a qualitative and quantitative comparison between MVTE and our multi-aspect embedding approach.

3. Bundle-based Queries

In this section, we describe our design for collecting crowdsource annotations in bundles which share a common aspect. In the following section, we describe how these bundled qualitative queries are embedded in corresponding multiple quantitative metric spaces.

The design of annotation task is the key for gathering information from a crowd. For a set of images $\mathcal{S} = \{O_i\}$, a typical task for acquiring similarity annotations is in this form: is image O_i more similar to O_j or O_k? The answer to the question is an ordered triplet: $T(i, j, k)$, in which O_i is more similar to O_j than O_k. Such triplet queries do not scale well to handle large datasets. Moreover, they do not provide context to the crowd worker, which can make the query more difficult to answer, and reduces the consistency of answers among multiple queries. It is possible that each query is answered with a different aspect in mind.

Each triplet query has to be associated with the correct embedding space to produce consistent embeddings. Thus, it is beneficial to use queries which gather similarity information for many triplets using the same aspect. To this end, inspired by [14], we ask workers to perform a clustering query. More specifically, a set of images is presented to each worker, who is requested to classify the images into multiple groups. Note that this classification task is cost-effective since it is only slightly harder than $T(i, j, k)$-like queries, but it can yield a large amount of $T(i, j, k)$-like annotations from these clusters.

More importantly, such a classification naturally leads the worker to use a single similarity aspect while performing the task. Thus, all the derived $T(i, j, k)$ triplets can be assigned to the same embedding space, i.e., they are bundled. The bundle queries greatly reduce the amount of affiliations to be inferred, as instead of inferring the affiliation of each triplet, only a single affiliation for each bundle is required.

Formally, in a query, we ask a crowd worker to classify $\mathbb{N}$ images into at most $\mathbb{B}$ bins/clusters $\{\mathcal{S}_c\}$. The aforementioned $T(i, j, k)$-like triplets can be derived from the clusters as $\{T(i, j, k)\}$, where, $O_i, O_j \in \mathcal{S}_x, O_k \in \mathcal{S}_y$ and $x \neq y$, i.e., two images from the same cluster are considered to be more similar than the third one from another cluster.

In practice, we chose to use a simpler representation of qualitative similarities — $T(i, j, \theta)$-like tuples. These tuples can be derived from the clusters $\{\mathcal{S}_c\}$ by producing a tuple $\{T(i, j, 1)\}$ where $O_i, O_j \in \mathcal{S}_x$ and a tuple $\{T(i, j, 0)\}$ where $O_i \in \mathcal{S}_x, O_j \in \mathcal{S}_y$. In other words, two images are considered to be similar/dissimilar if they are from same/different clusters. We denote tuples derived from query $q \in \mathcal{Q}$ as $\mathcal{T}^q = \{T(i, j, \theta)\}$. In the next section, we present a bundle optimization algorithm that takes bundled tuples $T(i, j, \theta)$ as input.

4. Multi-aspect Embedding in Bundles

As discussed above, a multitude of aspects cannot be captured in a consistent way within a single embedding space. Thus, we compute multiple embedding spaces $\mathcal{E} = \{E_s\}$ dedicated to different similarity aspects.

To associate bundled $T(i, j, \theta)$-tuples with appropriate embedding spaces, there are two sets of variables to solve. One set contains the aspect inference variables α_s^q, which indicates the likelihood that query q is based on the s-th similarity aspect. The other set contains the coordinates of the images in each embedding.

Let us denote the coordinates of image O_* in embedding E_s as $O_{*,s}$. We use constrastive loss [7] to model how well tuple $T(i, j, \theta)$ fits in the s-th embedding E_s:

$$\begin{aligned} L(T(i,j,\theta), E_s) &= \theta \times d(O_{i,s}, O_{j,s})^2 \\ &+(1-\theta) \times \max(0, m - d(O_{i,s}, O_{j,s}))^2, \end{aligned} \tag{1}$$

where $d(O_{i,s}, O_{j,s}) = \|O_{i,s} - O_{j,s}\|_2$ and m is a margin for embedding dissimilar images apart from each other. The loss of associating bundled tuples $\mathcal{T}^q$ with embedding E_s is then:

$$L(\mathcal{T}^q, E_s) = \sum_{T(i,j,\theta) \in \mathcal{T}^q} L(T(i, j, \theta), E_s). \tag{2}$$

Intuitively, $L(\mathcal{T}^q, E_s)$ is small when the tuples $\mathcal{T}^q$ from query q are associated with the embedding space that corresponds to the similarity aspect used by query q. However, it is unknown which embedding space is the best fit. We introduce aspect inference variables α_s^q to address this problem. Formally, the aggregate loss of bundled tuples $\mathcal{T}^q$ with respect to multiple embeddings $\{E_s\}$ is:

$$L(\mathcal{T}^q) = \sum_{E_s \in \mathcal{E}} \alpha_s^q \times L(\mathcal{T}^q, E_s), \quad \sum_{s=1}^{|\mathcal{E}|} \alpha_s^q = 1, \alpha_s^q > 0. \tag{3}$$

An inference variable α_s^q can be interpreted as the probability that query q is based on the s-th similarity aspect. As the optimization progresses, α_s^q gradually converge to associate query q with a specific embedding.

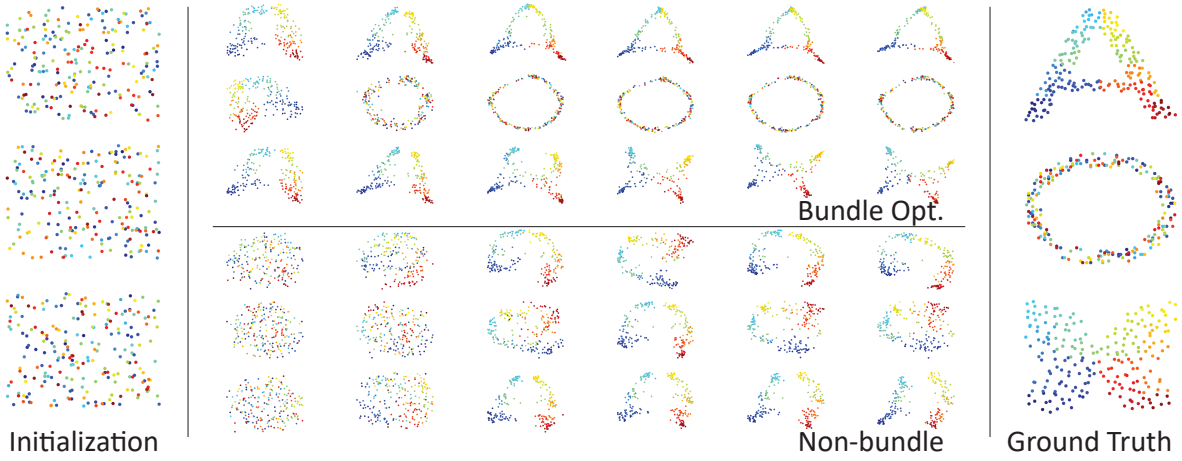

Figure 3: A visual comparison of embedding results w/wo bundle optimization at iteration 5, 10, 20, 30, 60 and 100.

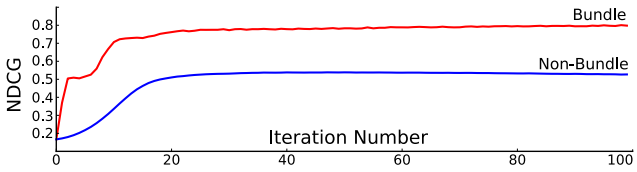

Figure 4: A comparison of Fig. 3 results in NDCG metric.

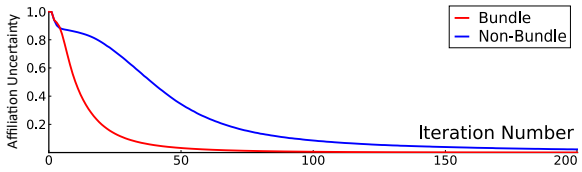

Figure 5: Affiliation uncertainty curve during recovery of "AO" embeddings, with and without bundle optimization.

Finally, we summarize the loss for all queries, and the optimization problem can be written as:

$$\arg\min_{\alpha_s^q,\ O_{*,s}} \sum_{q\in\mathcal{Q}} \sum_{E_s\in\mathcal{E}} \alpha_s^q \times \sum_{T(i,j,\theta)\in\mathcal{T}^q} L(T(i,j,\theta), E_s), \quad (4)$$

where $\sum_s \alpha_s^q = 1$, and $\alpha_s^q > 0$. Note that there is one aspect inference variable α_s^q per query per embedding, i.e., their total number is $|\mathcal{Q}| \times |\mathcal{E}|$. The loss function is differentiable with respect to variables α_s^q and $O_{*,s}$, thus it can be optimized with gradient descent based optimizers. We solve these two sets of variables simultaneously.

Initialization. If the embedding spaces are initialized with the same coordinates, or symmetrically with respect to the queries, the gradients are exactly the same. Thus, the gradient descent optimization updates them in the same way, which leads to identical embeddings. To avoid this, we use random initialization for the embedding coordinates. It can be assumed that the initial random embeddings are not equivalent or symmetric to one another with respect to the queries. However, in the beginning, such non-equivalence/asymmetry is probably weak, i.e., there is no strong tendency for a query to belong to a specific embedding. The non-equivalence/asymmetry is gradually reinforced by our algorithm, and the embeddings evolve into quite different spaces corresponding to multi-aspects.

We initialize α_s^q to $\frac{1}{|\mathcal{E}|}$, indicating that the queries have the same probability to be based on any of the unknown aspects. For specific applications, where relevant prior information can be leveraged, they can also be initialized with bias for different aspects.

5. Algorithm Analysis

Two distinctive features of our algorithm are the bundle optimization and the use of aspect inference variables. We study their effectiveness by extensive comparisons.

Analysis Settings. The goal of our algorithm is to estimate semantic similarity among images based on human perception. This similarity measure can be used as ground truth for supervised machine learning algorithms. However, the evaluation of our algorithm cannot be based on crowdsourcing data which is inexact in nature. Instead, we use synthetic data as ground truth to analyze and evaluate our algorithm.

We introduce a synthetic "AOB" dataset, which contains 214 points ($|\mathcal{S}| = 214$), distributed to form "A", "O", and "B"utterfly shapes in the ground truth embeddings $\mathcal{E}_{gt} = \{E_{gtA}, E_{gtO}, E_{gtB}\}$ (see Figure 3). The points are indexed sequentially according to point coordinates in E_{gtA} and E_{gtB}, so the embeddings are different but not completely

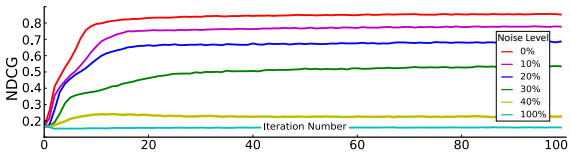

Figure 6: NDCG curve during the recovery of "AO" embeddings, with various levels of noise.

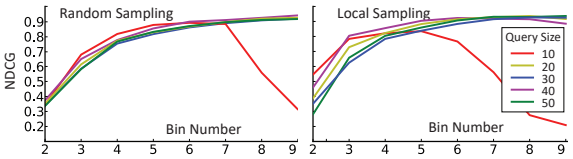

Figure 7: NDCG curve of "AO" embeddings recovery, using different query sampling, query size and bin number.

independent. In E_{gtO}, the points are indexed randomly, so that the embedding is completely independent of the other two. This way, the dataset simulates both dependent and independent aspects. We color the points by smoothly mapping their indices into continuous colors, *i.e.*, points with neighboring indices have similar colors.

We generate 600 random queries from each ground truth embedding and attempt to recover them by simulated query answers. Each query contains $\mathbb{N}(= 20)$ randomly sampled objects. Note that the random sampling strategy does not use any prior knowledge of the embeddings, to simulate actual crowdsourcing scenario where the ground truth is unknown. The answers are generated using K-means clustering of the samples, with $\mathbb{B}(= 5)$ seeds. The clustering is based on the position of objects in one of the embeddings (selected in random), to simulate users query answers which are based on a single unknown aspect. $114,000$ tuples are inferred from the clustering query answers of each embedding.

We evaluate the quality of recovered embeddings $\mathcal{E}$ based on the Normalized Discounted Cumulative Gain (NDCG) metric [12]. NDCG is widely used in evaluating retrieval relevance, and suitable for evaluating the recovery quality of the similarity based embedding spaces. More specifically, we first compute K-nearest ($K = 0.1 \times |\mathcal{S}|$) neighbors for each point in each recovered embedding in $\mathcal{E}$ and its corresponding embedding* in $\mathcal{E}_{gt}$, then with the corresponding ranked lists, NDCG are computed and averaged over all points (please refer to Appendix I for more details on NDCG calculation).

5.1. Bundle Optimization

A common approach for computing multi-embeddings is to collect and optimize over individual tuples. In this approach, each tuple may be based on a different aspect, thus the optimization has to infer aspects for each tuple individually. Formally, the non-bundled optimization problem can be written as:

$$\underset{\alpha_s^{(i,j,\theta)},\ O_{*,s}}{\arg\min} \sum_{E_s \in \mathcal{E}} \sum_{T(i,j,\theta) \in \mathcal{T}} \alpha_s^{(i,j,\theta)} \times L(T(i,j,\theta), E_s), \tag{5}$$

where $\alpha_s^{(i,j,\theta)}$ is the aspect inference variable indicating the probability of associating tuple $T(i,j,\theta) \in \mathcal{T}$ with embedding E_s, and $\sum_s \alpha_s^{(i,j,\theta)} = 1$, and $\alpha_s^{(i,j,\theta)} > 0$. The tuples can either be collected using single-tuple queries or inferred from a clustering query. Clearly, this formulation leads to many more variables to optimize than the bundled optimization, since the number of aspect inference variables is proportional to the number of tuples.

We apply bundled and non-bundled optimizations on AOB dataset, and show a visual comparison in Figure 3. Note that in both optimizations the embeddings are computed from the same random initialization. As can be seen in the figure, bundled optimization leads to a significantly better recovery of the ground truth embeddings than the non-bundled version. The bundled optimization (top row) produces distinct embeddings that resemble the ground truth, while the non-bundled optimization (bottom row) produces noisy embeddings that are quite similar to each other. The recovery quality is also evident in the color coding of the results. A high quality recovery should present color coding which is similar to the ground truth. While this is true for the bundled optimization results, in the non-bundled optimization results the color coding of all embeddings is similar only to the "A" and butterfly shapes and not to the "O" shape. This suggests that the aspects are not separated correctly, as all embeddings are influenced by tuples that represent similar color coding.

We also quantitatively measure the quality of the multi-aspect embedding recovery. At each iteration, we compute the average NDCG of the three embeding results, with bundled and non-bundled optimization. As can be seen in Figure 4, the bundled optimization converges much faster to more accurate embeddings.

5.2. Aspect Inference Variables

As discussed above, the optimization starts with random initialization of multiple embedding spaces, and progressively evolves to differentiate between distinct aspects. We examine this phenomenon in the task of recovering E_{gtA}

*In case of poor recovery, the correspondence between $\mathcal{E}$ and $\mathcal{E}_{gt}$ is not clear (see Figure 3 middle lower). In this case, we compute NDCG for all possible mappings between them, and pick the one with the highest NDCG as the most likely mapping.

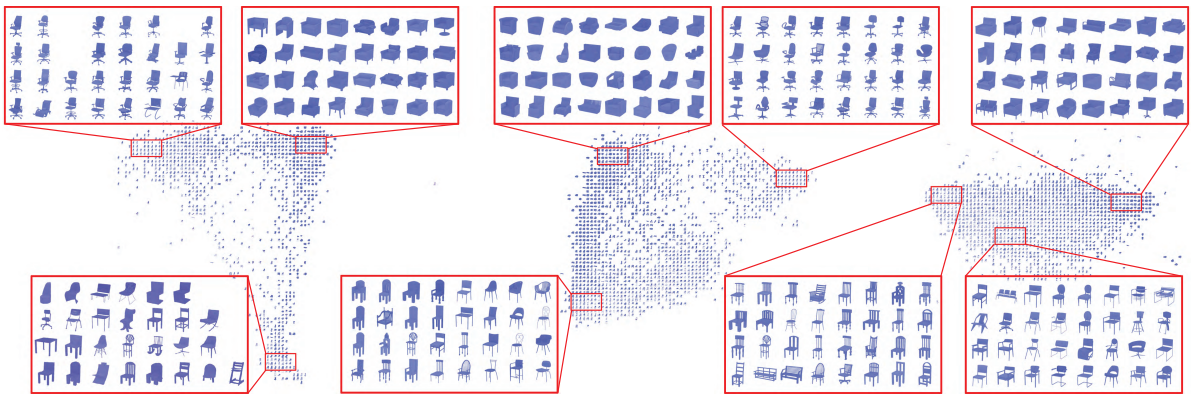

Figure 8: Recovering multiple aspects using our approach on the chair dataset with three aspects. The results show that the embeddings reflect specific aspects, from left to right: arms, legs and back of and chairs.

and E_{gtO} from simulated queries and answers, with an affiliation uncertainty metric:

$$\mathscr{U} = \frac{1}{|\mathcal{Q}|} * \sum_{q \in \mathcal{Q}} \frac{\min(\alpha_1^q, \alpha_2^q)}{\max(\alpha_1^q, \alpha_2^q)}, \tag{6}$$

where α_1^q and α_2^q are the aspect inference variables associated with E_A and E_O in query q. As shown in Figure 5, in the beginning, $\mathscr{U}$ is high, as the two initial embeddings are still in chaos state and it is not significant whether a query is associated with E_A or E_O. However, since E_A and E_O are not likely to be symmetric with respect to the queries, the asymmetry is gradually reinforced while one of the embeddings is evolving towards E_A and the other one towards E_O. This can be observed from the reduction of affiliation uncertainty as the optimization progresses.

Similarly, we can define the affiliation uncertainty metric for non-bundled optimization as:

$$\mathscr{U}_n = \frac{1}{|\mathcal{T}|} * \sum_{T(i,j,\theta) \in \mathcal{T}} \frac{\min(\alpha_1^{(i,j,\theta)}, \alpha_2^{(i,j,\theta)})}{\max(\alpha_1^{(i,j,\theta)}, \alpha_2^{(i,j,\theta)})}, \tag{7}$$

where $\alpha_1^{(i,j,\theta)}$ and $\alpha_2^{(i,j,\theta)}$ are the aspect inference variables associated with E_A and E_O for tuple $T(i,j,\theta)$. $\mathscr{U}_n$ is also plotted in Figure 5. $\mathscr{U}_n$ also reduces as the optimization progresses, which shows the aspect inference variables are somewhat effective without the bundled optimization. Still, the bundled optimization reduces affiliation uncertainty more effectively than the non-bundled version.

5.3. Stress Test under Noise

To test the robustness of our method to noise and erroneous measures we performed the following test. We collected the $228,000$ tuples from the clustering query answers of "A" and "O" embeddings, which form ground truth tuples. Then we add noise to the ground truth tuples by randomly selecting tuples and reversing their similar/dissimilar label. The portion of reversed tuples defines the noise level.

Figure 6 summarizes the results, where we can observe that our method tolerates small amount of noise, and the performance degrades smoothly. This indicates that the bundled optimization method is robust to reasonable degree of noisy measures.

5.4. Query Sampling, Query Size, and Bin Number

There are three factors controlling how we present each query to users: query sampling, query size $\mathbb{N}$ and bin number $\mathbb{B}$. Query sampling controls how the images clustered are sampled from the entire set, either randomly or locally (as proposed in [14]). In the local sampling strategy, the recovery is divided into several phases. Each phase first updates the embedings, and then sample points to form queries from neighbors, rather than randomly, in the following phase. $\mathbb{N}$ is the number of images to be clustered in each query. $\mathbb{B}$ is the number of bins — the maximum number of clusters allowed for each query. We evaluate the performance of our bundle optimization with synthetic data generated with different sampling strategy, $\mathbb{N}$ and $\mathbb{B}$.

We summarize the results in Figure 7. Note that, to make fair comparisons, we make sure $|\mathcal{Q}| \times \mathbb{N} = 3000$, i.e., more queries were used in tests with smaller query size. We can see that larger $\mathbb{N}$ and $\mathbb{B}$ give better recovery results, which is not surprising, since more information can be collected from such queries. On the other hand, smaller $\mathbb{N}$ and $\mathbb{B}$ are more friendly to crowd workers. In particular, a small value of $\mathbb{B}$ encourages workers to make decisions based on a single aspect in each query. We also show that the local sampling strategy proposed in [14], when extended to multiple-aspect setting, is effective with small query sizes and bin numbers. There exists a sweet point which balances

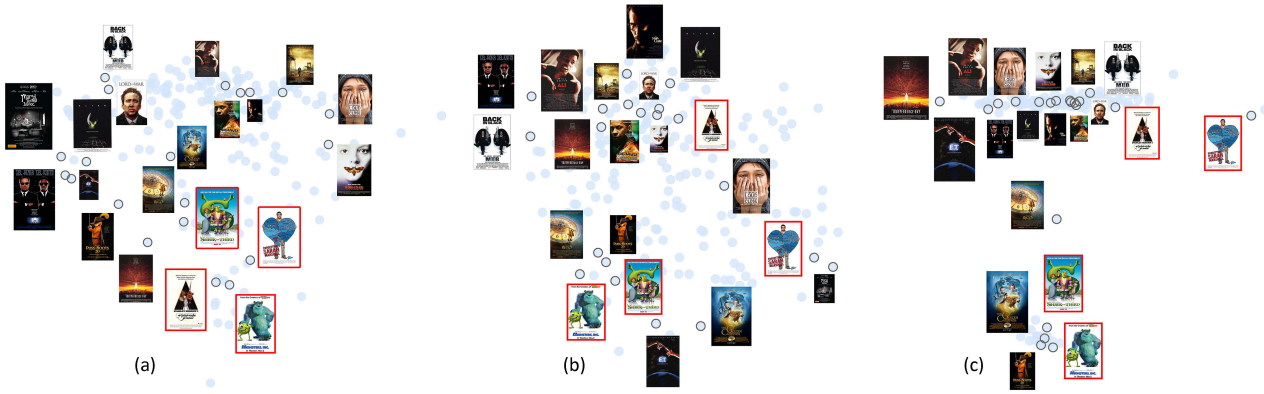

Figure 9: Recovering multiple aspects using our approach on the poster dataset with unknown aspects. The results show that the embeddings are ordered under different attributes: color and appearance (a), genre (b), animation / live action (c).

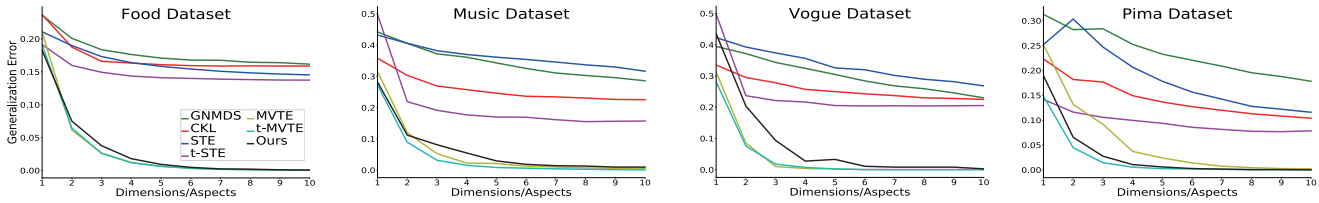

Figure 10: Comparison with previous single and multiple embedding methods on Food, Music, Vogue and Pima datasets.

these factors and maximize the cost-effectiveness.

6. Experimental Results

In this section we discuss our experimental results with data crowdsourcd from Amazon Mechanical Turk (AMT). We compare our results with relevant previous work for learning single or multiple aspects. We provide a dataset with clustering labels of three known aspects and a dataset with clustering labels of unknown aspects, both of which are collected from crowd workers via AMT.

Our optimization is implemented with Tensorflow [1]. In Equation 4 and 5, instead of directly optimizing α_s to satisfy $\sum_s \alpha_s = 1$ and $\alpha_s^q > 0$, we let $\alpha_s = softmax(\beta_s)$, and optimize β_s without constraints. We use Adam method [13] with learning rate 0.01 for the optimizations. All code and data will be opensourced.

6.1. AMT-based Experiments

In the following experiments we use AMT crowdsourced data to learn multi-aspect embeddings.

Recovering predefined aspects. To properly evaluate the performance of our method, we require a dataset that is intuitive enough for crowd queries, but also relates to known predefined aspects. To this end, we use a chair dataset which contains $6,777$ images rendered from the chair category of ShapeNet [5] and collect user data which related to multiple aspects in this dataset. In our AMT experiment, crowd workers were required to cluster queries considering one of the following predefined aspects: arms, legs and back of the chairs (please refer to Appendix II for AMT implementation and experimental details). We then recover appropriate embeddings of these aspects *without* the prior knowledge of the aspect the users were asked to consider. Figure 8 shows the recovered embeddings for three aspects. It can be seen that each embedding reflects a separate aspect. This demonstrates the ability of our method to automatically distinguish and classify the query answers according to their unknown aspects.

Recovering unknown aspects. In this experiment, we use our multi-aspect embedding method to embed movie poster images [14] using AMT queries which do not have any predefined aspect. That is, the crowd workers are not guided to use specific aspects, but can choose any aspect they see fit when clustering the images. We then compute three-aspect embeddings using our method. The results can be seen in Figure 9. For the purpose of display clarity, we only present 20 sampled images out of the total in the dataset. The nature of images in the dataset and the unconstrained settings of the experiment suggest the images may be categorized by the crowd workers using a large number of aspects. Thus, each embedding may reflect a mix of several aspects. Still, we can identify a meaningful

	AO	AOB	3D-2	3D-3
MVTE	0.47	0.49	0.28	0.31
t-MVTE	0.91	0.68	0.38	0.41
Ours	0.97	0.82	0.79	0.75

Table 1: NDCG performance compared with (t-)MVTE algorithms on various synthetic datasets.

distinction between the three embeddings. Embedding (a) reflects the appearance of the posters, in terms of color, composition and the content of the poster. For example, posters with white background (see marked images) appear close to each other in this emebedding, even though the movies belong to various genres. Embedding (b) reflects external context such as the genre of the movie or the actors that play in it. Note that horror movies appear on the top right, sci-fi movies appear on the top left, and family movies appear on the bottom. In embedding (c), the distinction between animated movies and live action movies takes precedence over other aspects, creating two tight groups of movies, animated and non-animated.

6.2. Quantitative Comparisons

We conduct an experiment to evaluate the effectiveness of our algorithm and compare it with (t-)MVTE multi-view embedding algorithm [3]. We use four synthetic datasets with multiple aspects, as well as different dimensions, as can be seen in Table 2. For example, 3D-2 is a set of three-dimensional points composed of two geometric models that correspond to two aspect distributions. Triplets and bundled tuples are generated from the ground truth data for each aspect and mixed together as input to recover multi-aspect embedding. Table 2 summarizes the NDCG of the recovered multi-embedding compared with the ground truth. Our algorithm outperforms (t-)MVTE, and is more stable when dealing with complex data of higher dimensions.

We also compare our method with several previous single embedding methods, including GNMDS [2], CKL [21], and (t-)STE [22] using datasets from previous work: Food [25], Music [8], Vogue [11] and Pima Indians Diabetes [20] dataset. The previous three datasets offer individually collected $\{T(i,j,k)\}$ similarity triplets. Note that Pima dataset is composed of ground truth numeric attributes[†], the triplets are generated by the method from [3], rather than crowdsourced. We convert each $\{T(i,j,k)\}$ triplet to $\{T(i,j,1)\}$ and $\{T(i,k,0)\}$ tuples, which are bundled together. Our algorithm degrades to our non-bundle baseline on these datasets, since only this minimal bundle information can be leveraged.

As in [3], generalization error is used to evaluate the performances, which describes the dissatisfaction ratio of new recovered triplets in the ground truth. For multi-embedding methods, a triplet is considered to be satisfied if its distance relationship in one of the multi-aspect embedding is consistent with that in ground truth. Since single embedding methods cannot recover multi-aspect embedding, we compare with their ability to recover corresponding high dimensional space. Figure 10 shows that multi-aspect embedding methods outperform single embedding methods, and the tendency is more obvious as dimensions/aspects increase. For these datasets, which do not contain bundled queries, our generalization errors are comparable to (t-)MVTE. Note that the expected performance using bundled queries in similar settings is quite higher, as discussed in Section 5.

Limitations. Different similarity aspects may have different popularity. Similarly to (t)-MVTE, our method does not take this into account. In addition, the number of different embeddings in our method needs to be manually set. While larger number of embeddings can better reflect more aspects, there is a risk that they actually represent noise or outlier measures. An interesting future work is to try and differentiate inliers and outliers measures, and automatically pick a suitable number of embedding spaces.

7. Conclusions

We have presented a method for multi-aspect embedding. The method takes qualitative measures and solves an optimization problem that embeds them into multiple spaces such that the quantitative measures in the embeded spaces agree with the qualitative measures. The optimization solves two sets of unknown parameters simultaneously: one is the embedded coordinates of the points and the other is the classification variables of the measure to the unknown aspects. We presented a bundle optimization and showed its power to infer the aspect classification. We showed that it outperforms existing multiple embedding methods. Our experimental results on crowdsourced data demonstrate the competence of our method to produce multi-embedding from inconsistent and redundant data.

Acknowledgements

Our work was supported by National Key Research & Development Plan of China (No. 2016YFB1001404) and National Natural Science Foundation of China (No. 61602273).

[†]Each element in Pima dataset is associated with three numeric attributes: plasma glucose concentration, diastolic blood pressure, and 2-hour serum insulin, which are the same as that in [3].

References

[1] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org. 7

[2] S. Agarwal, J. Wills, L. Cayton, G. Lanckriet, D. Kriegman, and S. Belongie. Generalized non-metric multidimensional scaling. In *Proc. of Int. Conf. on AI and Statistics*, San Juan, Puerto Rico, 2007. 2, 8

[3] E. Amid and A. Ukkonen. Multiview triplet embedding: Learning attributes in multiple maps. In *Proc.of ICML*, pages 1472C–1480, 2015. 2, 8, 10

[4] S. Bell and K. Bala. Learning visual similarity for product design with convolutional neural networks. *ACM TOG*, 34(4):98:1–98:10, July 2015. 2

[5] A. X. Chang, T. A. Funkhouser, L. J. Guibas, P. Hanrahan, Q. Huang, Z. Li, S. Savarese, M. Savva, S. Song, H. Su, J. Xiao, L. Yi, and F. Yu. Shapenet: An information-rich 3d model repository. *CoRR*, abs/1512.03012, 2015. 7, 10

[6] A. I. Chittilappilly, L. Chen, and S. Amer-Yahia. A survey of general-purpose crowdsourcing techniques. *IEEE TKDE*, 28(9):2246–2266, Sept 2016. 2

[7] S. Chopra, R. Hadsell, and Y. LeCun. Learning a similarity metric discriminatively, with application to face verification. In *Proc. of CVPR*, volume 1, pages 539–546. IEEE, 2005. 3

[8] D. P. Ellis, B. Whitman, A. Berenzweig, and S. Lawrence. The quest for ground truth in musical artist similarity. In *ISMIR*. Paris, France, 2002. 8

[9] E. Garces, A. Agarwala, D. Gutierrez, and A. Hertzmann. A similarity measure for illustration style. *ACM TOG*, 33(4):93:1–93:9, July 2014. 2

[10] A. Globerson and S. Roweis. Metric learning by collapsing classes. In *Proc. of NIPS*, pages 451–458, Cambridge, MA, USA, 2005. MIT Press. 2

[11] H. Heikinheimo and A. Ukkonen. The crowd-median algorithm. In *Proceedings of the First AAAI Conference on Human Computation and Crowdsourcing, HCOMP 2013, November 7-9, 2013, Palm Springs, CA, USA*, 2013. 8

[12] K. Järvelin and J. Kekäläinen. IR evaluation methods for retrieving highly relevant documents. In *Proc. of SIGIR*, pages 41–48, New York, NY, USA, 2000. ACM. 5, 10

[13] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *Proceedings of the 3rd International Conference on Learning Representations (ICLR)*, 2014. 7

[14] Y. Kleiman, G. Goldberg, Y. Amsterdamer, and D. Cohen-Or. Toward semantic image similarity from crowdsourced clustering. *Visual Computer*, 32(6-8):1045–1055, June 2016. 1, 2, 3, 6, 7, 10, 11

[15] J. Li and D. Jurafsky. Do multi-sense embeddings improve natural language understanding? In *Conf. on Empirical Methods for Natural Language Processing*, 2015. 2

[16] Y. Liu, Z. Liu, T.-S. Chua, and M. Sun. Topical word embeddings. In *Proc. of AAAI*, pages 2418–2424. AAAI Press, 2015. 2

[17] P. O'Donovan, J. Lībeks, A. Agarwala, and A. Hertzmann. Exploratory font selection using crowdsourced attributes. *ACM TOG*, 33(4):92:1–92:9, July 2014. 2

[18] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. Imagenet large scale visual recognition challenge. *Int. J. Comput. Vision*, 115(3):211–252, Dec. 2015. 2

[19] K. Rvelin, Kek, and J. Inen. Cumulated gain-based evaluation of IR techniques. *Acm Transactions on Information Systems*, 20(4):422–446, 2002. 10

[20] J. W. Smith, J. Everhart, W. Dickson, W. Knowler, and R. Johannes. Using the adap learning algorithm to forecast the onset of diabetes mellitus. In *Proc. of the Annual Symposium on Computer Application in Medical Care*, page 261. American Medical Informatics Association, 1988. 8

[21] O. Tamuz, C. Liu, S. Belongie, O. Shamir, and A. Kalai. Adaptively learning the crowd kernel. In L. Getoor and T. Scheffer, editors, *Proc. of ICML*, pages 673–680, New York, NY, USA, June 2011. ACM. 1, 2, 8

[22] L. van der Maaten and K. Weinberger. Stochastic triplet embedding. In *IEEE Int. Workshop on Machine Learning for Signal Processing*, pages 1–6, Sept 2012. 8

[23] J. Wang, H. Do, A. Woznica, and A. Kalousis. Metric learning with multiple kernels. In *Proc. of NIPS*, pages 1170–1178, USA, 2011. Curran Associates Inc. 2

[24] J. Wang, T. Kraska, M. J. Franklin, and J. Feng. Crowder: Crowdsourcing entity resolution. *Proc. VLDB Endow.*, 5(11):1483–1494, July 2012. 2

[25] M. J. Wilber, I. S. Kwak, and S. J. Belongie. Cost-effective hits for relative similarity comparisons. In *2nd AAAI Conference on Human Computation and Crowdsourcing*, 2014. 8

[26] Z. Wu and C. L. Giles. Sense-aware semantic analysis: A multi-prototype word representation model using wikipedia. In *Proc. of AAAI*, pages 2188–2194. AAAI Press, 2015. 2

[27] P. Xie and E. P. Xing. Multi-modal distance metric learning. In *Proc. of IJCAI*, pages 1806–1812. AAAI Press, 2013. 2

[28] E. P. Xing, A. Y. Ng, M. I. Jordan, and S. Russell. Distance metric learning, with application to clustering with side-information. In *Proc. of NIPS*, pages 521–528, Cambridge, MA, USA, 2002. MIT Press. 2

A. Appendix I: Comparison Details

We hereby present a comprehensive version of our multi-embedding results, as well as those generated by previous work, and more details of the implementation of quantitative comparisons.

NDCG Calculation. We evaluate the quality of multi-embedding based on Normalized Discounted Cumulative Gain (NDCG) metric [12, 19]. NDCG is widely used for evaluating information retrieval, and is suitable for evaluating the recovery quality of an embedding space. The NDCG for a point p in an embedding space is defined as $\text{NDCG}_p = \frac{\text{DCG}_p}{\text{IDCG}_p}$, and the NDCG of an embedding space is defined as the average NDCG of the points in the space. The term DCG_p is the Discounted Cumulative Gain for point p, and is defined as:

$$\text{DCG}_p = \sum_{i=1}^{K} \frac{2^{rel_i} - 1}{log_2(i+1))}, \tag{8}$$

where $K = 0.1 \times |\mathcal{S}|$ is the number of nearest neighbors in the embedding space to be evaluated, and rel_i is the relevance between p and its i-th nearest neighbor p_i. The relevance is defined as:

$$rel_i = e^{-\frac{d(p,p_i)}{d(p,p_K)}}, \tag{9}$$

where $d(p, p_i)$ denotes the distance between p and p_i in the *ground truth* embedding space. The term IDCG_p is the ideal DCG_p, *i.e.*, DCG_p computed in the ground truth embedding space. For a perfectly recovered embedding space, DCG_p is the same as $IDCG_p$, producing $NDCG_p = 1$. A higher NDCG value indicates a better approximation of the ground truth embedding.

Comparisons with (t-)MVTE. MVTE [3] is the most relevant previous work to our method, as both share the same goal of recovering multi-aspect embeddings. To compare the performance of producing multi-aspect embeddings using both methods, we conducted a comprehensive comparison with (t-)MVTE on synthetic dataset both qualitatively and quantitatively. Please note that t-MVTE is an implementation of MVTE with a different distribution assumption. The results of (t-)MVTE were produced using the authors' open sourced implementation.

Table 2 summarizes the qualitative and quantitative (NDCG) comparisons between our method and (t-)MVTE on the four synthetic datasets (AO, AOB, 3D-2 and 3D-3) described in the paper. Our method produces higher quality embeddings in all test cases.

The numeric attributes of the elements in Pima dataset form the three ground truth embedding spaces on three aspects. With these ground truth embedding spaces, we can generate synthetic queries, recover multiple embeddings from the queries and evaluate them in terms of NDCG score, same as those done for the synthetic datasets. The NDCG evaluation is summarized in the last column of Table 2. It is clear that our method with bundle information outperforms (t-)MVTE on Pima dataset, which is consistent with the comparisons observed on the synthetic datasets.

B. Appendix II: AMT Data Collections

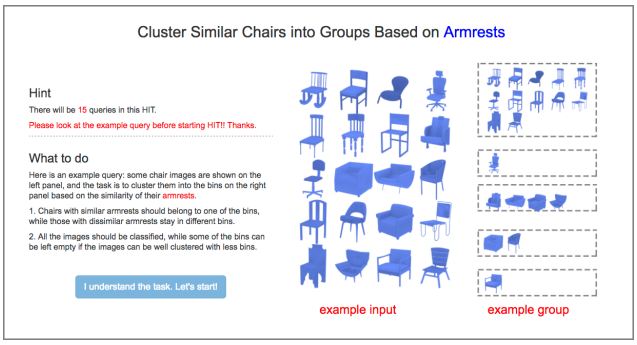
Cluster Similar Chairs into Groups Based on Armrests
Hint
There will be 15 queries in this HIT.
Please look at the example query before starting HIT!! Thanks.
What to do
Here is an example query: some chair images are shown on the left panel, and the task is to cluster them into the bins on the right panel based on the similarity of their armrests.
1. Chairs with similar armrests should belong to one of the bins, while those with dissimilar armrests stay in different bins.
2. All the images should be classified, while some of the bins can be left empty if the images can be well clustered with less bins.
I understand the task. Let's start!
example input
example group

Figure 11: MTurk HIT introduction for predefined-aspect (*e.g.*, arm) experiment on Chair dataset.

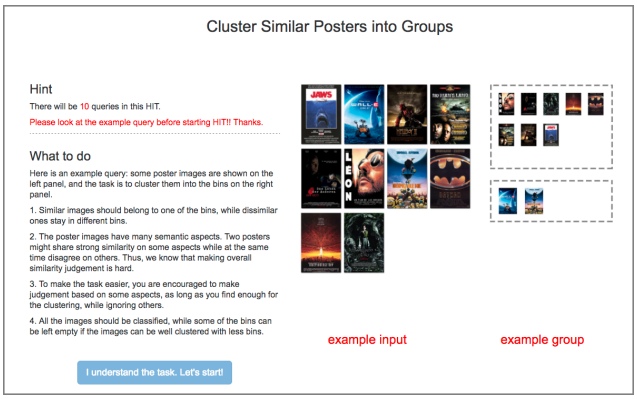
Cluster Similar Posters into Groups
Hint
There will be 10 queries in this HIT.
Please look at the example query before starting HIT!! Thanks.
What to do
Here is an example query: some poster images are shown on the left panel, and the task is to cluster them into the bins on the right panel.
1. Similar images should belong to one of the bins, while dissimilar ones stay in different bins.
2. The poster images have many semantic aspects. Two posters might share strong similarity on some aspects while at the same time disagree on others. Thus, we know that making overall similarity judgement is hard.
3. To make the task easier, you are encouraged to make judgement based on some aspects, as long as you find enough for the clustering, while ignoring others.
4. All the images should be classified, while some of the bins can be left empty if the images can be well clustered with less bins.
I understand the task. Let's start!
example input
example group

Figure 12: MTurk HIT introduction for unknown-aspect experiment on Poster dataset.

To properly evaluate our performance on crowd answers which relate to both known and unknown aspects, we designed human subject experiments in Amazon Mechanical Turk to cluster semantic similarity queries. We use a chair dataset with $6,777$ images rendered from the chair category of ShapeNet [5], and a poster dataset with 200 randomly selected film posters from [14]. In the chair dataset experiment, crowd workers were asked to cluster queries considering one of the following *predefined* aspects: arms, legs and back of the chairs. In the poster dataset experiment, crowd workers were asked to cluster queries considering *any* aspects that they prefer.

	AO	AOB	3D-2	3D-3	Pima
			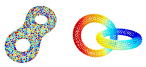		
MVTE	0.47	0.49	0.28	0.31	0.34
t-MVTE	0.91	0.68	0.38	0.41	0.59
	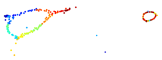		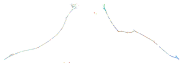		
Ours	0.97	0.82	0.79	0.75	0.75
	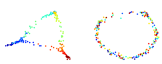		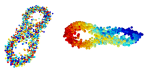		

Table 2: NDCG performance and embedding visualization compared with (t-)MVTE algorithms on various datasets.

Predefined-aspects Experiment on Chair Dataset. We collected semantic similarity data using clustering queries as in [14] instead of the more traditional triplet queries, using a drag-and-drop graphical UI with 20 clustering images shown on the left and 5 grouping bins on the right. Workers were required to cluster the 20 images into the bins, with similar ones in the same bin. In Amazon Mechanical Turk, an experiment for each worker is considered as a human intelligence task (HIT). In our experiment, each HIT begins with an examplar introduction with guidelines for crowd workers (see Figure 11), followed by 15 queries. As a quality control, 3 queries in each HIT are ground truth sentinels, answers from a worker with lower than 70% sentinel accuracy will be rejected. Additionally, only crowd workers with higher than 80% approval rate can accept our HIT.

We distribute in total 45,000 queries with ground truth sentinels for the predefined-aspects experiment, producing 3,000 HITs. After quality control, we collected 41,287 valid clustering query answers, which are aggregated to 7,953,827 final tuples. Workers are paid on average $0.25 to cluster 15 queries and spent an average of 6 minutes per HIT. The total cost was about $800 spent over roughly a month and a half.

Unknown-aspects Experiment on Poster Dataset. Film posters present rich semantic information, which makes it hard for workers to cluster queries without any predefined aspects. To simplify the task, we collect clustering queries using a drag-and-drop graphical UI with 10 clustering images shown on the left and 2 grouping bins on the right. Workers were asked to cluster the 10 images into the bins according to their own preferences. In this experiment, each HIT begins with an examplar introduction with guidelines for crowd workers (see Figure 12), followed by 10 queries. Since there is no ground truth answers for this experiment, we constrain valid crowd workers as those higher than 80% approval rate as a quality control.

In total, we distributed 840 queries, producing in total 84 HITs. After excluding incomplete answers, we collected 800 valid queries which are aggregated to 36,000 final tuples. Workers are paid on average $0.25 to cluster 10 queries and spent an average of 3 minutes per response. Please note that the reward was higher in this experiment to accelerate the data collection. The total cost was about $24 spent over one day.